arXiv:1702.03561v1 [math.NA] 12 Feb 2017

Modern Monte Carlo Variants for Uncertainty Quantification in Neutron Transport

Ivan G. Graham, Matthew J. Parkinson, and Robert Scheichl

Abstract We describe modern variants of Monte Carlo methods for Uncertainty Quantification (UQ) of the Neutron Transport Equation, when it is approximated by the discrete ordinates method with diamond differencing. We focus on the mono-energetic 1D slab geometry problem, with isotropic scattering, where the cross-sections are log-normal correlated random fields of possibly low regularity. The paper includes an outline of novel theoretical results on the convergence of the discrete scheme, in the cases of both spatially variable and random cross-sections. We also describe the theory and practice of algorithms for quantifying the uncertainty of a linear functional of the scalar flux, using Monte Carlo and quasi-Monte Carlo methods, and their multilevel variants. A hybrid iterative/direct solver for computing each realisation of the functional is also presented. Numerical experiments show the effectiveness of the hybrid solver and the gains that are possible through quasi-Monte Carlo sampling and multilevel variance reduction. For the multilevel quasi-Monte Carlo method, we observe gains in the computational ε-cost of up to 2 orders of magnitude over the standard Monte Carlo method, and we explain this theoretically. Experiments on problems with up to several thousand stochastic dimensions are included.

Keywords: Reactor Modelling, Neutron (Boltzmann) Transport Equation, Monte Carlo, QMC, MLMC, Source Iteration.

1 Introduction

In this paper we will consider the Neutron Transport equation (NTE), sometimes referred to as the Boltzmann transport equation. This is an integro-differential equation which models the flux of neutrons in a reactor. It has particular applications for

Ivan G. Graham · Matthew J. Parkinson · Robert Scheichl (✉)
University of Bath, Claverton Down Road, Bath, UK
e-mail: i.g.graham@bath.ac.uk; m.parkinson@bath.ac.uk; r.scheichl@bath.ac.uk

nuclear reactor design, radiation shielding and astrophysics [43]. There are many potential sources of uncertainty in a nuclear reactor, such as the geometry, material composition and reactor wear. Here, we will consider the problem of random spatial variation in the coefficients (the *cross-sections*) in the NTE, represented by correlated random fields with potentially low smoothness. Our aim is to understand how *uncertainty in the cross-sections propagates through to (functionals of) the neutron flux*. This is the forward problem of Uncertainty Quantification.

We will quantify the uncertainty using Monte Carlo (MC) type methods, that is, by simulating a finite number of pseudo-random instances of the NTE and by averaging the outcome of those simulations to obtain statistics of quantities of interest. Each statistic can be interpreted as an expected value of some (possibly nonlinear) functional of the neutron flux with respect to the random cross-sections. The input random fields typically need to be parametrised with a significant number of random parameters leading to a problem of high-dimensional integration. MC methods are known to be particularly well-suited to this type of problem due to their dimension independent convergence rates.

However, convergence of the MC algorithm is slow and determined by $\sqrt{\mathbb{V}(\cdot)/N}$, where $\mathbb{V}(\cdot)$ is the variance of the quantity of interest and N is the number of samples. For this reason, research is focussed on improving the convergence, whilst retaining dimensional independence. Advances in MC methods can broadly be split into two main categories: improved sampling and variance reduction. Improved sampling methods attempt to find samples that perform better than the pseudo-random choice. Effectively, they aim to improve the $\sqrt{1/N}$ term in the error estimate. A major advance in sampling methods has come through the development of quasi-Monte Carlo (QMC) methods. Variance reduction methods, on the other hand, attempt to reduce the $\mathbb{V}(\cdot)$ term in the error estimate and thus reduce the number of samples needed for a desired accuracy. Multilevel Monte Carlo (MLMC) methods (initiated in [28, 18] and further developed in, e.g., [20, 7, 10, 9, 33, 31, 46, 27]) fall into this category. A comprehensive review of MLMC can be found in [19].

The rigorous theory of all of the improvements outlined above requires regularity properties of the solution, the verification of which can be a substantial task. There are a significant number of published papers on the regularity of parametric elliptic PDEs, in physical and parameter space, as they arise, e.g., in flow in random models of porous media [9, 32, 12, 13, 24, 33, 31]. However, for the NTE, this regularity question is almost untouched. Our complementary paper [25] contains a full regularity and error analysis of the discrete scheme for the NTE with spatially variable and random coefficients. Here we restrict to a summary of those results.

The field of UQ has grown very quickly in recent years and its application to neutron transport theory is currently of considerable interest. There are a number of groups that already work on this problem, e.g. [4, 17, 21] and references therein. Up to now, research has focussed on using the polynomial chaos expansion (PCE), which comes in two forms; the intrusive and non-intrusive approaches. Both approaches expand the random flux in a weighted sum of orthogonal polynomials. The intrusive approach considers the expansion directly in the differential equation, which in turn requires a new solver ('intruding' on the original solver). In contrast,

the non-intrusive approach attempts to estimate the coefficients of the PCE directly, by projecting onto the PCE basis cf. [4, eq.(40)]. This means the original solver can be used as a 'black box' as in MC methods. Both of the approaches then use quadrature to estimate the coefficients in the PCE. The main disadvantage of standard PCE is that typically the number of terms grow exponentially in the number of stochastic dimensions and in the order of the PCE, the so-called *curse of dimensionality*.

Fichtl and Prinja [17] were some of the first to numerically tackle the 1D slab geometry problem with random cross-sections. Gilli et al. [21] improved upon this work by using (adaptive) sparse grid ideas in the collocation method, to tackle the curse of dimensionality. Moreover, [5] constructed a hybrid PCE using a combination of Hermite and Legendre polynomials, observing superior convergence in comparison to the PCE with just Hermite polynomials. More recently [4] tackled the (time-independent) full criticality problem in three spatial, two angular and one energy variable. They consider a second expansion, the high-dimensional model representation (HDMR), which allows them to expand the response (e.g. functionals of the flux) in terms of low-dimensional subspaces of the stochastic variable. The PCE is used on the HDMR terms, each with their own basis and coefficients. We note however, that none of these papers provide any rigorous error or cost analysis.

The structure of this paper is as follows. In Section 2, we describe the model problem, a 1D slab geometry simplification of the Neutron Transport Equation with spatially varying and random cross-sections. We set out the discretisation of this equation and discuss two methods for solving the resultant linear systems; a direct and an iterative solver. In Section 3, the basic elements of a fully-discrete error analysis of the discrete ordinates method with diamond differencing applied to the model problem are summarised. The full analysis will be given in [25]. In Section 4, we introduce a number of variations on the Monte Carlo method for quantifying uncertainty. This includes a summary of the theoretical computational costs for each method. Finally, Section 5 contains numerical results relating to the rest of the paper. We first present a hybrid solver that combines the benefits of both direct and iterative solvers. Its cost depends on the particular realisation of the cross-sections. Moreover, we present simulations for the UQ problem for the different variants of the Monte Carlo methods, and compare the rates with those given by the theory.

2 The Model Problem

The Neutron Transport Equation (NTE) is a physically derived balance equation, that models the angular flux $\psi(\mathbf{r},\Theta,E)$ of neutrons in a domain, where $\mathbf{r}$ is position, Θ is angle and E is energy. Neutrons are modelled as non-interacting particles travelling along straight line paths with some energy E. They interact with the larger nuclei via absorption, scattering and fission. The rates σ_A, σ_S and σ_F at which these events occur are called the *absorption, scattering and fission cross-sections*, respectively. They can depend on the position $\mathbf{r}$ and the energy E of the neutron. The

scattering cross-sections also depend on the energy E' after the scattering event, as well as on the angles Θ and Θ' before and after the event.

The two main scenarios of interest in neutron transport are the so-called *fixed source problem* and the *criticality problem*. We will focus on the former, which concerns the transport of neutrons emanating from some fixed source term f. It has particular applications in radiation shielding. We will further simplify our model to the *1D slab geometry case* by assuming

- no energy dependence;
- dependence only on one spatial dimension and infinite extent of the domain in the other two dimensions;
- no dependence of any cross-sections on angle;
- no fission.

The resulting simplified model is an integro-differential equation for the angular flux $\psi(x,\mu)$ such that

$$\mu \frac{\mathrm{d}\psi}{\mathrm{d}x}(x,\mu) \;+\; \sigma(x)\psi(x,\mu) \;=\; \sigma_S(x)\phi(x) \;+\; f(x)\,, \tag{1}$$

$$\text{where} \qquad \phi(x) \;=\; \frac{1}{2}\int_{-1}^{1} \psi(x,\mu')\,\mathrm{d}\mu'\,, \tag{2}$$

for any $x \in (0,1)$ and $\mu \in [-1,1]$, subject to the no in-flow boundary conditions

$$\psi(0,\mu) \;=\; 0, \;\text{ for }\; \mu > 0 \quad \text{and} \quad \psi(1,\mu) = 0, \;\text{ for }\; \mu < 0\,. \tag{3}$$

Here, the angular domain is reduced from $\mathbb{S}_2$ to the unit circle $\mathbb{S}_1$ and parametrised by the cosine $\mu \in [-1,1]$ of the angle. The equation degenerates at $\mu = 0$, i.e. for neutrons moving perpendicular to the x-direction. The coefficient function $\sigma(x)$ is the total cross-section given by $\sigma = \sigma_S + \sigma_A$. For more discussion on the NTE see [11, 36].

2.1 Uncertainty Quantification

An important problem in industry is to quantify the uncertainty in the fluxes due to uncertainties in the cross-sections. Most materials, in particular shielding materials such as concrete, are naturally heterogeneous or change their properties over time through wear. Moreover, the values of the cross-sections are taken from nuclear data libraries across the world and they can differ significantly between libraries [35]. This means there are large amounts of uncertainty on the coefficients, and this could have significant consequences on the system itself.

To describe the random model, let $(\Omega, \mathscr{A}, \mathbb{P})$ be a probability space with $\omega \in \Omega$ denoting a random event from this space. Consider a (finite) set of partitions of the spatial domain, where on each subinterval we assume that $\sigma_S = \sigma_S(x,\omega)$ and $\sigma = \sigma(x,\omega)$ are two (possibly dependent or correlated) random fields. Then the

angular flux and the scalar flux become random fields and the model problem (1), (2) becomes

$$\mu \frac{\mathrm{d}\psi}{\mathrm{d}x}(x,\mu,\omega) \;+\; \sigma(x,\omega)\psi(x,\mu,\omega) \;=\; \sigma_S(x,\omega)\phi(x,\omega) \;+\; f(x)\;, \tag{4}$$

$$\text{where} \qquad \phi(x,\omega) \;=\; \int_{-1}^{1} \psi(x,\mu',\omega)d\mu' \tag{5}$$

and $\psi(\cdot,\cdot,\omega)$ satisfies the boundary conditions (3). The set of equations (4), (5), (3) have to hold for almost all realisations $\omega \in \Omega$.

For simplicity, we restrict ourselves to deterministic $\sigma_A = \sigma_A(x)$ with

$$0 \;<\; \sigma_{A,\min} \;\leq\; \sigma_A(x) \;\leq\; \sigma_{A,\max} \;<\; \infty\;, \quad \text{for all} \quad x \in [0,1]\;, \tag{6}$$

and assume a log-normal distribution for $\sigma_S(x,\omega)$. The total cross-section $\sigma(x,\omega)$ is then simply the log-normal random field with values $\sigma(x,\omega) = \sigma_S(x,\omega) + \sigma_A(x)$. In particular, we assume that $\log\sigma_S$ is a correlated zero mean Gaussian random field, with covariance function defined by

$$C_\nu(x,y) \;=\; \sigma_{var}^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(2\sqrt{\nu}\frac{|x-y|}{\lambda_C}\right)^{\nu} K_\nu \left(2\sqrt{\nu}\frac{|x-y|}{\lambda_C}\right)\;. \tag{7}$$

This class of covariances is called the Matérn class. It is parametrised by the smoothness parameter $\nu \geq 0.5$; λ_C is the correlation length, σ_{var}^2 is the variance, Γ is the gamma function and K_ν is the modified Bessel function of the second kind. The limiting case, i.e. $\nu \to \infty$, corresponds to the Gaussian covariance function $C_\infty(x,y) \;=\; \sigma_{var}^2 \exp(-|x-y|^2/\lambda_C^2)$.

To sample from σ_S we use the Karhunen-Loève (KL) expansion of $\log\sigma_S$, i.e.,

$$\log\sigma_S(x,\omega) \;=\; \sum_{i=1}^{\infty} \sqrt{\xi_i}\;\eta_i(x)\;Z_i(\omega)\;, \tag{8}$$

where $Z_i \sim \mathcal{N}(0,1)$ i.i.d. Here ξ_i and η_i are the eigenvalues and the $L^2(0,1)$-orthogonal eigenfunctions of the covariance integral operator associated with kernel given by the covariance function in (7). In practice, the KL expansion needs to be truncated after a finite number of terms (here denoted d). The accuracy of this truncation depends on the decay of the eigenvalues [37]. For $\nu < \infty$, this decay is algebraic and depends on the smoothness parameter ν. In the Gaussian covariance case the decay is exponential. Note that for the Matérn covariance with $\nu = 0.5$, the eigenvalues and eigenfunctions can be computed analytically [37]. For other cases of ν, we numerically compute the eigensystem using the Nyström method - see, for example, [16].

The goal of stochastic uncertainty quantification is to understand how the randomness in σ_S and σ propagates to functionals of the scalar or angular flux. Such quantities of interest may be point values, integrals or norms of ϕ or ψ. They are random variables and the focus is on estimating their mean, variance or distribution.

2.2 Discretisation

For each realisation $\omega \in \Omega$, the stochastic 1D NTE (4), (5), (3) is an integro-differential equation in two variables, space and angle. For ease of presentation, we suppress the dependency on $\omega \in \Omega$ for the moment.

We use a $2N$-point quadrature rule $\int_{-1}^{1} f(\mu)d\mu \approx \sum_{|k|=1}^{N} w_k f(\mu_k)$ with nodes $\mu_k \in [-1,1]\backslash\{0\}$ and positive weights w_k to discretise in angle, assuming the (anti-) symmetry properties $\mu_{-k} = -\mu_k$ and $w_{-k} = w_k$. (In later sections, we construct such a rule by using N-point Gauss-Legendre rules on each of $[-1,0)$ and $(0,1]$.)

To discretise in space, we introduce a mesh $0 = x_0 < x_1 < \ldots < x_M = 1$ which is assumed to resolve any discontinuities in the cross-sections σ, σ_S and is also quasiuniform - i.e. the subinterval lengths $h_j := x_j - x_{j-1}$ satisfy $\gamma h \le h_j \le h := \max_{j=1,\ldots M} h_j$, for some constant $\gamma > 0$. Employing a simple Crank-Nicolson method for the transport part of (4), (5) and combining it with the angular quadrature rule above we obtain the classical *diamond-differencing* scheme:

$$\mu_k \frac{\Psi_{k,j} - \Psi_{k,j-1}}{h} \; + \; \sigma_{j-1/2} \frac{\Psi_{k,j} + \Psi_{k,j-1}}{2} \\ = \; \sigma_{S,j-1/2} \Phi_{j-1/2} \; + \; F_{j-1/2} \; , \quad j = 1,...,M, \;\; |k| = 1,\ldots,N, \tag{9}$$

where

$$\Phi_{j-1/2} \; = \; \frac{1}{2} \sum_{|k|=1}^{N} w_k \frac{\Psi_{k,j} \; + \; \Psi_{k,j-1}}{2} \; , \;\; j = 1,...,M \; . \tag{10}$$

Here $\sigma_{j-1/2}$ denotes the value of σ at the mid-point of the interval $I_j = (x_{j-1}, x_j)$, with the analogous meaning for $\sigma_{S,j-1/2}$ and $F_{j-1/2}$. The notation reflects the fact that (in the next section) we will associate the unknowns $\Psi_{k,j}$ in (9) with the nodal values $\psi_{k,h}(x_j)$ of continuous piecewise-linear functions $\psi_{k,h} \approx \psi(\cdot, \mu_k)$.

Finally, (9) and (10) have to be supplemented with the boundary conditions $\Psi_{k,0} = 0$, for $k > 0$ and $\Psi_{k,M} = 0$, for $k < 0$. If the right-hand side of (9) were known, then (9) could be solved simply by sweeping from left to right (when $k > 0$) and from right to left (when $k < 0$). The appearance of $\Phi_{j-1/2}$ on the right-hand side means that (9) and (10) consitute a coupled system with solution $(\Psi, \Phi) \in \mathbb{R}^{2NM} \times \mathbb{R}^M$. It is helpful to think of Ψ as being composed of $2N$ subvectors Ψ_k, each with M entries $\Psi_{k,j}$, consisting of approximations to $\psi(x_j, \mu_k)$ with x_j ranging over all free nodes.

The coupled system (9) and (10) can be written in matrix form as

$$\begin{pmatrix} T & -\Sigma_S \\ -P & I \end{pmatrix} \begin{pmatrix} \Psi \\ \Phi \end{pmatrix} \; = \; \begin{pmatrix} F \\ \mathbf{0} \end{pmatrix} \; . \tag{11}$$

Here, the vector $\Phi \in \mathbb{R}^M$ contains the approximations of the scalar flux at the M midpoints of the spatial mesh. The matrix T is a block diagonal $2NM \times 2NM$ matrix, representing the left hand side of (9). The $2N$ diagonal blocks of T, one per angle, are themselves bi-diagonal. The $2NM \times M$ matrix Σ_S simply consists of $2N$ identical diagonal blocks, one per angle, representing the multiplication of Φ by σ_S at the

midpoints of the mesh. The $M \times 2NM$ matrix P represents the right hand side of (10), i.e. averaging at the midpoints and quadrature. The matrix I denotes the $M \times M$ identity matrix. The vector $F \in \mathbb{R}^{2NM}$ contains $2N$ copies of the source term evaluated at the M midpoints of the spatial mesh.

2.3 Direct and Iterative Solvers

We now wish to find the (approximate) fluxes in the linear system (11). We note that the matrix T is invertible and has a useful sparsity structure that allows its inverse to be calculated in $\mathcal{O}(MN)$ operations. However, the bordered system (11) is not as easy to invert, due to the presence of Σ_S and P.

To exploit the sparsity of T, we do block elimination on (11) obtaining the Schur complement system for the scalar flux, i.e.,

$$\left(I - PT^{-1}\Sigma_S\right)\Phi \;=\; PT^{-1}F \;, \tag{12}$$

which now requires the inversion of a smaller (dense) matrix. Note that (12) is a finite-dimensional version of the reduction of the integro-differential equation (4), (5) to the integral form of the NTE, see (20). In this case, the two dominant computations with $\mathcal{O}(M^2N)$ and $\mathcal{O}(M^3)$ operations respectively, are the triple matrix product $PT^{-1}\Sigma_S$ in the construction of the Schur complement and the LU factorisation of the $M \times M$ matrix $\left(I - PT^{-1}\Sigma_S\right)$. This leads to a total

$$\text{theoretical cost of the direct solver } \sim\; \mathcal{O}(M^2(M\,+\,N))\;. \tag{13}$$

We note that for stability reasons (see §3, also [41] in a simpler context), the number of spatial and angular points should be related. A suitable choice is $M \sim N$, leading to a cost of the direct solver of $\mathcal{O}(M^3)$ in general.

The second approach for solving (11) is an iterative solver commonly referred to as *source iteration*, cf. [8]. The form of (12) naturally suggests the iteration

$$\Phi^{(k)} \;=\; PT^{-1}\left(\Sigma_S\Phi^{(k-1)} \,+\, F\right)\;, \tag{14}$$

where $\Phi^{(k)}$ is the approximation at the kth iteration, with $\Phi^{(0)} = PT^{-1}F$. This can be seen as a discrete version of an iterative method for the integral equation (20).

In practice, we truncate after K iterations. The dominant computations in the source iteration are the K multiplications with $PT^{-1}\Sigma_S$. Exploiting the sparsity of all the matrices involved, these multiplications cost $\mathcal{O}(MN)$ operations, leading to an overall

$$\text{theoretical cost of source iteration } \sim\; \mathcal{O}\left(MNK\right)\;. \tag{15}$$

Our numerical experiments in Section 5 show that for $N = 2M$ the hidden constants in the two estimates (13) and (15) are approximately the same. Hence, whether the

iterative solver is faster than the direct solver depends on whether the number of iterations K to obtain an accurate enough solution is smaller or larger than M.

There are sharp theoretical results on the convergence of source iteration for piecewise smooth cross-sections [8, Thm 2.20]. In particular, if $\phi^{(K)}(\omega)$ denotes the approximation to $\phi(\omega)$ after K iterations, then

$$\left\| \sigma^{1/2} \left(\phi - \phi^{(K)} \right) \right\|_2 \leq C' \left(\eta \left\| \frac{\sigma_S}{\sigma} \right\|_\infty \right)^K , \tag{16}$$

for some constant C' and $\eta \leq 1$. That is, the error decays geometrically with rate no slower than the spatial maximum of σ_S/σ. This value depends on ω and will change pathwise. Using this result as a guide together with (6), we assume that the convergence of the L^2-error with respect to K can be bounded by

$$\|\phi - \phi^{(K)}\|_2 \leq C \left\| \frac{\sigma_S}{\sigma} \right\|_\infty^K , \tag{17}$$

for some constant C that we will estimate numerically in Section 5.

3 Summary of Theoretical Results

The rigorous analysis of UQ for PDEs with random coefficients requires estimates for the error when discretisations in physical space (e.g. by finite elements) and probability space (e.g. by sampling techniques) are combined. The physical error estimates typically need to be probabilistic in form (e.g. estimates of expectation of the physical error). Such estimates are quite well-developed for elliptic PDEs - see for example [9] but this question is almost untouched for the transport equation (or more specifically the NTE). We report here (somewhat cautiously) some preliminary results which will be presented in detail in the forthcoming paper [25]. Our procedure is to first give an error analysis for (1), (2) with variable cross-sections, which is explicit in σ, σ_S and then to use this to give probabilistic error estimates for the spatial discretisation (9), (10).

The numerical analysis of the NTE (and related integro-differential equation problems such as radiative transfer) dates back at least as far as the work of H.B. Keller [29]. After a huge growth in the mathematics literature in the 1970's and 1980's, progress has been slower since. This is perhaps surprising, since discontinuous Galerkin (DG) methods have enjoyed a massive recent renaissance and the neutron transport problem was one of the motivations behind the original introduction of DG [42]. Even today, an error analysis of the NTE with variable (even deterministic) cross-sections (with explicit dependence on the data) is still not available, even for the model case of mono-energetic 1D slab geometry considered here.

The fundamental paper on the analysis of the discrete ordinates method for the NTE is [41]. Here a full analysis of the combined effect of angular and spatial dis-

cretisation is given under the assumption that the cross-sections σ and σ_S in (4) are constant. The delicate relation between spatial and angular discretisation parameters required to achieve stability and convergence is described there. Later research e.g. [2], [3] produced analogous results for models of increasing complexity and in higher dimensions, but the proofs were mostly confined to the case of cross-sections that are constant in space. A separate and related sequence of papers (e.g. [34], [47], and [1]) allow for variation in cross-sections, but error estimates explicit in this data are not available there.

Our results here are orientated to the case when σ, σ_S have relatively rough fluctuations. As a precursor to the random case introduced in §2, we assume that there is some partition of $[0,1]$ and that σ, σ_S are C^α functions on each subinterval of the partition (with $\alpha \in (0,1]$), but that σ, σ_S may be discontinuous across the break points. We assume that the mesh x_j introduced in §2.2 resolves these break points. (Here C^α is the usual Hölder space of index α with norm $\|\cdot\|_\alpha$.) We also assume that the source function $f \in C^\alpha$.

When discussing the error when (9), (10) is applied to (1), (2), it is useful to consider the pure transport problem

$$\mu \frac{\mathrm{d}u}{\mathrm{d}x} + \sigma u = g, \quad \text{with } u(0) = 0, \text{ when } \mu > 0 \quad \text{and } u(1) = 0 \text{ when } \mu < 0, \quad (18)$$

and with $g \in C$ a generic right-hand side. Application of the Crank-Nicolson method (as in (9)) yields

$$\mu \left(\frac{U_j - U_{j-1}}{h_j} \right) + \sigma_{j-1/2} \left(\frac{U_j + U_{j-1}}{2} \right) = g_{j-1/2} \quad (19)$$

with analogous boundary conditions, where, for any continuous function c, we use $c_{j-1/2}$ to denote $c(x_{j-1/2})$. Letting V^h denote the space of continuous piecewise linear functions with respect to the mesh $\{x_j\}$, (19) is equivalent to seeking a $u^h \in V^h$ (with nodal values U_j) such that

$$\int_{I_j} \left(\mu \frac{\mathrm{d}u^h}{\mathrm{d}x} + \widetilde{\sigma} u^h \right) \; = \; \int_{I_j} \widetilde{g} \, , \quad \text{where} \quad I_j = (x_{j-1}, x_j) \, ,$$

and $\widetilde{c}$ denotes the piecewise constant function with respect to the grid $\{x_j\}$ which interpolates c at the mid-points of subintervals. It is easy to show that both (18) and (19) have unique solutions and we denote the respective solution operators by $\mathscr{S}_\mu$ and $\mathscr{S}^h_\mu$, i.e.

$$u = \mathscr{S}_\mu g \quad \text{and} \quad u^h = \mathscr{S}^h_\mu g \, .$$

Bearing in mind the angular averaging process in (2) and (10), it is useful to introduce the continuous and discrete spatial operators:

$$(\mathscr{K} g)(x) := \frac{1}{2} \int_{-1}^{1} \left(\mathscr{S}_\mu g \right)(x) \mathrm{d}\mu, \quad \text{and} \quad (\mathscr{K}^{h,N} g)(x) \; = \; \frac{1}{2} \sum_{|k|=1}^{N} w_k (\mathscr{S}^h_{\mu_k} g)(x) \, .$$

It is easy to see (and classically very well-known) that

$$(\mathscr{K}g)(x) = \frac{1}{2}\int_0^1 E_1(|\tau(x,y)|)g(y)\mathrm{d}y,$$

where E_1 is the exponential integral and $\tau(x,y) = \int_x^y \sigma$ is the optical path. In fact (even when σ is merely continuous), $\mathscr{K}$ is a compact Fredholm integral operator on a range of function spaces and $\mathscr{K}^{h,N}$ is a finite rank approximation to it. The study of these integral operators in the deterministic case is a classical topic, e.g. [44]. In the case of random σ, $\mathscr{K}$ is an integral operator with a random kernel which merits further investigation. Returning to (1), (2), we see readily that

$$\psi(x,\mu) = \mathscr{S}_\mu(\sigma_S\phi + f)\ , \quad \text{so that} \quad \phi \ = \ \mathscr{K}(\sigma_S\phi + f). \tag{20}$$

Moreover (9) and (10) correspond to a discrete analogue of (20) as follows. Introduce the family of functions $\psi_k^{h,N} \in V^h$, $|k| = 1,\ldots,N$, by requiring $\psi_k^{h,N}$ to have nodal values $\Psi_{k,j}$. Then set

$$\phi^{h,N} := \frac{1}{2}\sum_{|k|=1}^{N} w_k \psi_k^{h,N} \ \in V^h,$$

and it follows that (9) and (10) may be rewritten

$$\int_{I_j}\left(\mu_k\frac{\mathrm{d}\psi_k^{h,N}}{\mathrm{d}x} + \widetilde{\sigma}\psi_k^{h,N}\right) \ = \ \int_{I_j}\widetilde{g^{h,N}}\ , \quad \text{where} \quad g^{h,N} = \sigma_S\phi^{h,N} + f\ .$$

and thus

$$\psi_k^{h,N} = \mathscr{S}_{\mu_k}^h\left(\sigma_S\phi^{h,N} + f\right), \quad \text{so that} \quad \phi^{h,N} = \mathscr{K}^{h,N}(\sigma_S\phi^{h,N} + f)\ . \tag{21}$$

The numerical analysis of (9) and (10) is done by analysing (the second equation in) (21) as an approximation of the corresponding equation in (20). This is studied in detail in [41] for constant σ, σ_S. Here we consider the variable case, obtaining all estimates explicitly in σ, σ_S. Elementary manipulation on (20) and (21) shows that

$$\phi - \phi^{h,N} = (I - \mathscr{K}^{h,N}\sigma_S)^{-1}(\mathscr{K} - \mathscr{K}^{h,N})(\sigma_S\phi + f), \tag{22}$$

and so

$$\|\phi - \phi^{h,N}\|_\infty \le \|(I - \mathscr{K}^{h,N}\sigma_S)^{-1}\|_\infty \|(\mathscr{K} - \mathscr{K}^{h,N})(\sigma_S\phi + f)\|_\infty. \tag{23}$$

We then estimate each term on the right-hand side as follows:

(i) We show first that

$$\|(I - \mathscr{K}\sigma_S)^{-1}\|_\infty \ \lesssim \ \left(\frac{\sigma_{\max}}{\sigma_{\min}}\right)^{3/2}\left(\frac{1}{1 - \|\sigma_S/\sigma\|_\infty}\right)\ , \tag{24}$$

where the the symbol $\lesssim$ indicates that we suppress constants which are independent of σ_S, σ, h, or N. This estimate itself is not trivial, since for invertibility of $(I-\mathscr{K}\sigma_S)$ we need to appeal to the contractivity of $\mathscr{K}\sigma_S$ in a certain weighted L_2 space ([8]) and then convert the resulting estimate to the uniform norm. Details will be in [25].

(ii) We then show that $\|\mathscr{K}-\mathscr{K}^{h,N}\|_\infty$ is small as $h\to 0$ and $N\to\infty$.

Combining (ii) together with (i) and the Banach Lemma ensures that

$$\|(I-\mathscr{K}^{h,N}\sigma_S)^{-1}\|_\infty \le 2\|(I-\mathscr{K}\sigma_S)^{-1}\|_\infty, \tag{25}$$

for h sufficiently small and N sufficiently large. This establishes the stability of method (21). The actual values of h,N which are sufficient to ensure that the bound (25) holds depends on the cross-sections. Thus, unlike the elliptic case, the stability here is *path-dependent*. In order to make the later estimates digestible in this short summary, we combine (24) and (25) in the form

$$\|(I-\mathscr{K}^{h,N}\sigma_S)^{-1}\|_\infty \lesssim \mathscr{P}(\sigma,\sigma_S), \tag{26}$$

where (from now on) $\mathscr{P}(\sigma,\sigma_S)$ denotes a generic function of the cross sections of the form

$$(\sigma_{\max})^p(\sigma_{\min})^{-q}(1-\|\sigma/\sigma_S\|_\infty)^{-1}m(\|\sigma\|_\alpha)m(\|\sigma_S\|_\alpha)\,, \quad \text{with} \quad p,q\ge 0, \tag{27}$$

where $m(t):=\max\{1,t\}$. In the second part of the proof we estimate the second factor on the right-hand side of (23), i.e. the consistency error. Introducing the semidiscrete operator:

$$(\mathscr{K}^N f)(x) \;=\; \frac{1}{2}\sum_{|k|=1}^{N} w_k(\mathscr{S}_{\mu_k}f)(x)$$

(corresponding to applying the quadrature rule in angle but without discretisation in space), we then write $\mathscr{K}-\mathscr{K}^{h,N}=(\mathscr{K}-\mathscr{K}^N)+(\mathscr{K}^N-\mathscr{K}^{h,N})$. The consistency error is then estimated in two steps:

(iii) The semidiscrete error is

$$(\mathscr{K}-\mathscr{K}^N)(\sigma_S\phi+f)=\frac{1}{2}\left(\int_{-1}^{1}\psi(x,\mu)\mathrm{d}\mu-\sum_{|k|=1}^{N}w_k\psi(x,\mu_k)\right), \tag{28}$$

and the estimate for this needs an analysis of the regularity of ψ with respect to μ (explicit in the cross-sections).

(iv) The remaining component in the consistency estimate is

$$(\mathscr{K}^N-\mathscr{K}^{h,N})(\sigma_S\phi+f)=\frac{1}{2}\sum_{|k|=1}^{N}w_k\left(\mathscr{S}_{\mu_k}-\mathscr{S}^h_{\mu_k}\right)(\sigma_S\phi+f). \tag{29}$$

This is estimated by proving stability of the Crank-Nicolson method and a cross-section-explicit bound on $\|\phi\|_\alpha$. The end result is the following.

Theorem 1. *Under the assumptions outlined above,*

$$\|\phi - \phi^{h,N}\|_\infty \lesssim \left(\mathscr{P}_1 \, h \log N + \mathscr{P}_2 \, h^\alpha + \mathscr{P}_3 \, \frac{1}{N} \right) \|F\|_\alpha$$

with $\mathscr{P}_i = \mathscr{P}_i(\sigma, \sigma_S)$ being cross-section dependent constants of the form (27).

Returning to the case when σ, σ_S are random functions, this theorem provides pathwise estimates for the error. These can be turned into estimates in the corresponding Bochner space provided the coefficients $\mathscr{P}_i$ are bounded in probability space. This can then be decided once the random model for σ, σ_S is given. For the specific scenario outlined in §2 it can be shown, using the results in [9, §2], [24], that $\mathscr{P}_i \in L^p(\Omega)$ for all $1 \le p < \infty$, and hence we have:

Corollary 1. $\|\phi - \phi^{h,N}\|_{L^p(\Omega, L^\infty(0,1))} \lesssim \left(h \log N + h^\alpha + \frac{1}{N} \right) \|f\|_\alpha$.

4 Modern Variants of Monte Carlo

Let $Q(\omega) \in \mathbb{R}$ denote a functional of ϕ or ψ representing a quantity of interest. We will focus on estimating $\mathbb{E}[Q]$, the expected value of Q. Since we are not specific about what functionals we are considering, this includes also higher order moments or CDFs of quantities of interest. The expected value is a high-dimensional integral and the goal is to find efficient quadrature methods in high dimensions. We consider Monte Carlo type sampling methods.

As outlined above, to obtain samples of $Q(\omega)$ the NTE has to be approximated numerically. First, the random scattering cross section σ_S in (4) is sampled using the KL expansion of $\log \sigma_S$ in (8) truncated after d terms. The stochastic dimension d is chosen sufficiently high so that the truncation error is smaller than the other approximation errors. Let $Z^n \in \mathbb{R}^d$ be a realisation of the multivariate Gaussian coefficient $Z := (Z_i)_{i=1,\dots,d}$ in the KL expansion (8). For each $n \in \mathbb{N}$, let us denote by $Q_h(Z^n)$ the approximation of the nth sample of Q obtained numerically using a spatial grid with mesh size h and $2N$ angular quadrature points. We assume throughout that $N \sim 1/h$, so there is a single discretisation parameter h.

We will consider various unbiased, sample-based estimators $\widehat{Q}_h$ for the expected value $\mathbb{E}[Q]$ and we will quantify the accuracy of each estimator by its mean square error (MSE) $e(\widehat{Q}_h)^2$. Since $\widehat{Q}_h$ is assumed to be an unbiased estimate of $\mathbb{E}[Q_h]$, i.e. $\mathbb{E}[\widehat{Q}_h] = \mathbb{E}[Q_h]$, the MSE can be expanded as

$$e(\widehat{Q}_h)^2 \;=\; \mathbb{E}\left[(\widehat{Q}_h - \mathbb{E}[Q])^2 \right] \;=\; \left(\mathbb{E}\left[Q - Q_h \right] \right)^2 + \mathbb{V}[\widehat{Q}_h] \,, \tag{30}$$

i.e., the squared bias due to the numerical approximation plus the sampling (or quadrature) error $\mathbb{V}[\widehat{Q}_h] = \mathbb{E}[(\widehat{Q}_h - \mathbb{E}[Q_h])^2]$. In order to compare computational

costs of the various methods we will consider their ε-cost $\mathscr{C}_\varepsilon$, that is, the number of floating point operations to achieve a MSE $e(\widehat{Q}_h)^2$ less than ε^2.

To bound the ε-cost for each method, we make the following assumptions on the discretisation error and on the average cost to compute a sample from Q_h:

$$\left|\mathbb{E}[Q-Q_h]\right| \;=\; \mathscr{O}(h^\alpha)\;, \tag{31}$$

$$\mathbb{E}[\mathscr{C}(Q_h)] \;\leq\; \mathscr{O}(h^{-\gamma})\;, \tag{32}$$

for some constants $\alpha,\gamma>0$. We have seen in Section 2 that (32) holds with γ between 2 and 3. The new theoretical results in Section 3 guarantee that (31) also holds for some $\alpha>0$. However, the results in Section 3 are for general coefficients of of low smoothness. The practically observed values for α in our numerical experiments are significantly bigger with values between 1.5 and 2.

In recent years, many alternative methods for high-dimensional integrals have emerged that use tensor product deterministic quadrature rules combined with sparse grid techniques to reduce the computational cost [48, 6, 39, 26, 4, 17, 21]. The efficiency of these approaches relies on high levels of smoothness of the parameter to output map and in general their cost may grow exponentially with the number of parameters (the *curse of dimensionality*). Such methods are not competitive with Monte Carlo type methods for our problem, where the coefficients are parametrised by an infinite number of independent random variables and we typically need to keep a large number of these parameters for a reasonable accuracy.

However, standard Monte Carlo methods are notoriously slow to converge, requiring thousands or even millions of samples to achieve acceptable accuracies. In our application, where each sample involves the numerical solution of an integro-differential equation this very easily becomes intractable. The novel Monte Carlo approaches that we present here, aim to improve this situation in two complementary ways. Quasi-Monte Carlo methods reduce the number of samples to achieve a certain accuracy dramatically by using deterministic ideas to find well distributed samples in high dimensions. Multilevel methods use the available hierarchy of numerical approximations to our integro-differential equation to shift the bulk of the computations to cheap, inaccurate coarse models while providing the required accuracy with only a handful of expensive, accurate model solves.

4.1 Standard Monte Carlo

The (standard) Monte Carlo (MC) estimator for $\mathbb{E}[Q]$ is defined by

$$\widehat{Q}_h^{MC} \;:=\; \frac{1}{N_{MC}}\sum_{n=1}^{N_{MC}} Q_h(Z^n)\;, \tag{33}$$

where N_{MC} is the number of Monte Carlo points/samples $Z^n \sim \mathcal{N}(0,I)$. The sampling error of this estimator is $\mathbb{V}[\widehat{Q}_h^{MC}] = \mathbb{V}[Q_h]/N_{MC}$.

A sufficient condition for the MSE to be less than ε^2 is for both the squared bias and the sampling error in (30) to be less than $\varepsilon^2/2$. Due to assumption (31), a sufficient condition for the squared bias to be less than $\varepsilon^2/2$ is $h \sim \varepsilon^{1/\alpha}$. Since $\mathbb{V}[Q_h]$ is bounded with respect to $h \to 0$, the sampling error of $\widehat{Q}_h^{MC}$ is less than $\varepsilon^2/2$ for $N_{MC} \sim \varepsilon^{-2}$. With these choices of h and N_{MC}, it follows from Assumption (32) that the mean ε-cost of the standard Monte Carlo estimator is

$$\mathbb{E}\left[\mathscr{C}_\varepsilon(\widehat{Q}_h^{MC})\right] \;=\; \mathbb{E}\left[\sum_{n=1}^{N_{MC}} \mathscr{C}(Q_h(Z^n))\right] \;=\; N_{MC}\,\mathbb{E}[\mathscr{C}(Q_h)] \;=\; \mathscr{O}\left(\varepsilon^{-2-\frac{\gamma}{\alpha}}\right). \quad (34)$$

To reduce this high ε-cost is the goal of the novel Monte Carlo approaches which we will discuss now.

4.2 Quasi-Monte Carlo

The first approach to reduce the ε-cost is based on using quasi-Monte Carlo (QMC) rules, which replace the random samples in (33) by carefully chosen deterministic samples and treat the expected value with respect to the d-dimensional Gaussian Z in (8) as a high-dimensional integral with Gaussian measure.

Initially interest in QMC points arose within number theory in the 1950's, and the theory is still at the heart of good QMC point construction today. Nowadays, the fast component-by-component construction (CBC) [40] provides a quick method for generating good QMC points, in very high-dimensions. Further information on the best choices of deterministic points can be found, for example in [45, 15, 38, 14].

The choice of QMC points can be split into two categories; lattice rules and digital nets. We will only consider randomised rank-1 lattice rules here. In particular, given a suitable generating vector $z \in \mathbb{Z}^d$ and R independent, uniformly distributed random shifts $(\Delta_r)_{r=1}^R$ in $[0,1]^d$, we construct $N_{QMC} = RP$ lattice points in the unit cube $[0,1]^d$ using the simple formula

$$v^{(n)} = \text{frac}\left(\frac{nz}{P} + \Delta_r\right), \qquad n = 1,\dots,P,\;\; r = 1,\dots,R$$

where "frac" denotes the fractional part function applied componentwise. To transform the lattice points $v^n \in [0,1]^d$ into "samples" $\widetilde{Z}^n \in \mathbb{R}^d$, $n = 1,\dots,N_{QMC}$, of the multivariate Gaussian coefficients Z in the KL expansion (8) we apply the inverse cumulative normal distribution. See [23] for details.

Finally, the QMC estimator is given by

$$\widehat{Q}_h^{QMC} \;:=\; \frac{1}{N_{QMC}} \sum_{n=1}^{N_{QMC}} Q_h(\widetilde{Z}^n)\,,$$

Note that this is essentially identical in its form to the standard MC estimator (33), but crucially with deterministically chosen and then randomly shifted $\widetilde{Z}^n$. The random shifts ensure that the estimator is unbiased, i.e. $\mathbb{E}[\widehat{Q}_h^{QMC}] = \mathbb{E}[Q_h]$.

The bias for this estimator is identical to the MC case, leading again to a choice of $h \sim \varepsilon^{1/\alpha}$ to obtain a MSE of ε^2. Here the MSE corresponds to the mean square error of a randomised rank-1 lattice rule with P points averaged over the shift $\Delta \sim \mathscr{U}([0,1]^d)$. In many cases, it can be shown that the quadrature error, i.e., the second term in (30), converges with $\mathscr{O}(N_{QMC}^{-1/2\lambda})$, with $\lambda \in (\frac{1}{2}, 1]$. That is, we can potentially achieve $\mathscr{O}(N_{QMC}^{-1})$ convergence for $\widehat{Q}_h^{QMC}$ as opposed to the $\mathscr{O}(N_{MC}^{-1/2})$ convergence for $\widehat{Q}_h^{MC}$. The rate of convergence depends on the smoothness of the map $Z \to Q_h(Z)$ and the generating vector z does in theory have to be chosen problem specific. However, standard generating vectors, such as those available at [30], seem to also work well (and better than MC samples). Furthermore, we note the recent developments in "higher-order nets" [22, 12], which potentially increase the convergence of QMC methods to $\mathscr{O}(N_{QMC}^{-q})$, for $q \geq 2$.

Given the improved rate of convergence of the quadrature error and fixing the number of random shifts to $R = 8$, it suffices to choose $P \sim \varepsilon^{-2\lambda}$ for the quadrature error to be $\mathscr{O}(\varepsilon^2)$. Therefore it follows again from Assumption (32) that the ε-cost of the QMC method satisfies

$$\mathbb{E}_\Delta\left[\mathscr{C}_\varepsilon(\widehat{Q}^{QMC})\right] \;=\; \mathscr{O}\left(\varepsilon^{-2\lambda-\frac{\gamma}{\alpha}}\right) . \tag{35}$$

When $\lambda = \frac{1}{2}$ (or close to it), this is essentially a reduction in the ε-cost by a whole order of ε. In the case of non-smooth random fields, we typically have $\lambda \approx 1$ and the ε-cost grows with the same rate as that of the standard MC method. However, the absolute cost is always reduced, even in that case.

4.3 Multilevel Methods

The main issue with the above methods is the high cost for computing the samples $\{Q_h(Z^{(n)})\}$, each requiring us to solve the NTE. The idea of the multilevel Monte Carlo (MLMC) method is to use a hierarchy of discrete models of increasing cost and accuracy, corresponding to a sequence of decreasing discretisation parameters $h_0 > h_1 > \ldots > h_L = h$. Here, only the most accurate model on level L is designed to give a bias of $\mathscr{O}(\varepsilon)$ by choosing $h_L = h \sim \varepsilon^{1/\alpha}$ as above. The bias of the other models can be significantly higher.

MLMC methods were first proposed in an abstract way for high-dimensional quadrature by Heinrich [28] and then popularised in the context of stochastic differential equations in mathematical finance by Giles [18]. MLMC methods were first applied in uncertainty quantification in [7, 10]. The MLMC method has quickly gained popularity and has been further developed and applied in a variety of other problems. See [19] for a comprehensive review. In particular, the multilevel ap-

proach is not restricted to standard MC estimators and can also be used in conjunction with QMC estimators [20, 33, 31] or with stochastic collocation [46]. Here, we consider multilevel variants of standard MC and QMC.

MLMC methods exploit the linearity of the expectation, writing

$$\mathbb{E}[Q_h] \;=\; \sum_{\ell=0}^{L} \mathbb{E}[Y_\ell]\,, \qquad \text{where } \; Y_\ell := Q_{h_\ell} - Q_{h_{\ell-1}} \;\text{ and }\; Q_{h_{-1}} := 0.$$

Each of the expected values on the right hand side is then estimated separately. In particular, in the case of a standard MC estimator with N_ℓ samples for the ℓth term, we obtain the MLMC estimator

$$\widehat{Q}_h^{MLMC} \;:=\; \sum_{\ell=0}^{L} \widehat{Y}_\ell^{MC} \;=\; \sum_{\ell=0}^{L} \frac{1}{N_\ell} \sum_{n=1}^{N_\ell} Y_\ell(Z^{\ell,n})\,. \tag{36}$$

Here, $\{Z^{\ell,n}\}_{n=1}^{N_\ell}$ denotes the set of i.i.d. samples on level ℓ, chosen independently from the samples on the other levels.

The key idea in MLMC is to avoid estimating $\mathbb{E}[Q_h]$ directly. Instead, the expectation $\mathbb{E}[Y_0] = \mathbb{E}[Q_{h_0}]$ of a possibly strongly biased, but cheap approximation of Q_h is estimated. The bias of this coarse model is then estimated by a sum of correction terms $\mathbb{E}[Y_\ell]$ using increasingly accurate and expensive models. Due to the decreasing variance of Y_ℓ, as $h_\ell \to 0$, the number of samples N_ℓ to achieve a prescribed accuracy on level ℓ can be gradually reduced, leading to a lower overall cost of the MLMC estimator. More specifically, we have the following cost savings:

- On the coarsest level, using (32), the cost per sample is reduced from $\mathscr{O}(h^{-\gamma})$ to $\mathscr{O}(h_0^{-\gamma})$. Provided $\mathbb{V}[Q_{h_0}] \approx \mathbb{V}[Q_h]$ and h_0 can be chosen independently of ε, the cost of estimating $\mathbb{E}[Q_{h_0}]$ to an accuracy of ε in (36) is reduced to $\mathscr{O}(\varepsilon^{-2})$.
- On the finer levels, the number of samples N_ℓ to estimate $\mathbb{E}[Y_\ell]$ to an accuracy of ε in (36) is proportional to $\mathbb{V}[Y_\ell]\varepsilon^{-2}$. Now, provided $\mathbb{V}[Y_\ell] = \mathscr{O}(h_\ell^\beta)$, for some $\beta > 0$, which is guaranteed if Q_{h_ℓ} converges almost surely to Q pathwise, then we can reduce the number of samples as $h_\ell \to 0$. Depending on the actual values of α, β and γ, the cost to estimate $\mathbb{E}[Y_L]$ on the finest level can, in the best case, be reduced to $\mathscr{O}(\varepsilon^{-\gamma/\alpha})$.

The art of MLMC is to balance the number of samples across the levels to minimise the overall cost. This is a simple constrained optimisation problem to achieve $\mathbb{V}[\widehat{Q}_h^{MLMC}] \le \varepsilon^2/2$. As shown in [18], using the technique of Lagrange Multipliers, the optimal number of samples on level ℓ is given by

$$N_\ell \;=\; \left\lceil 2\varepsilon^{-2} \left(\sum_{\ell=0}^{L} \sqrt{\mathbb{V}[Y_\ell]/\mathscr{C}_\ell} \right) \sqrt{\mathbb{V}[Y_\ell]\mathscr{C}_\ell} \right\rceil\,, \tag{37}$$

where $\mathscr{C}_\ell := \mathbb{E}[\mathscr{C}(Y_\ell)]$. In practice, it is necessary to estimate $\mathbb{V}[Y_\ell]$ and $\mathscr{C}_\ell$ in (37) from the computed samples, updating N_ℓ as the simulation progresses.

Using these values of N_ℓ it is possible to establish the following theoretical complexity bound for MLMC [10].

Theorem 2. *Let us assume that* (31) *and* (32) *hold with* $\alpha, \gamma > 0$ *and that there exists* $\beta > 0$ *such that*

$$\mathbb{V}[Y_\ell] \;=\; \mathscr{O}(h_\ell^\beta)\,. \tag{38}$$

Then, with $L \sim \log(\varepsilon^{-1})$ *and with the choice of* $\{N_\ell\}_{l=0}^L$ *in* (37) *we have*

$$\mathbb{E}\left[\mathscr{C}_\varepsilon(\widehat{Q}_{h_L}^{MLMC})\right] \;=\; \mathscr{O}\Big(\varepsilon^{-2-\max\left(0,\ \frac{\gamma-\beta}{\alpha}\right)}\Big)\,. \tag{39}$$

When $\beta = \gamma$*, then there is an additional factor* $\log(\varepsilon^{-1})$.

Using lattice points $\widetilde{Z}^{\ell,n}$, as defined in Section 4.2, instead of the random samples $Z^{\ell,n}$ we can in the same way define a multilevel quasi-Monte Carlo (MLQMC) estimator

$$\widehat{Q}_h^{MLQMC} \;:=\; \sum_{\ell=0}^{L} \widehat{Y}_\ell^{QMC} \;=\; \sum_{\ell=0}^{L} \frac{1}{\widetilde{N}_\ell} \sum_{n=1}^{\widetilde{N}_\ell} Y_\ell(\widetilde{Z}^{\ell,n})\,. \tag{40}$$

The optimal values for $\widetilde{N}_\ell$ can be computed in a similar way to those in the MLMC method. However, they depend strongly on the rate of convergence of the lattice rule and in particular on the value of λ which is difficult to estimate accurately. We will give a practically more useful approach below.

It is again possible to establish a theoretical complexity bound, cf. [33, 31].

Theorem 3. *Let us assume that* (31) *and* (32) *hold with* $\alpha, \gamma > 0$ *and that there exists* $\lambda \in (\frac{1}{2}, 1]$ *and* $\beta > 0$ *such that*

$$\mathbb{V}_\Delta[\widehat{Y}_\ell^{QMC}] \;=\; \mathscr{O}\left(\widetilde{N}_\ell^{-1/\lambda} h_\ell^\beta\right)\,. \tag{41}$$

Let the number of random shifts on each level be fixed to R *and let* $L \sim \log(\varepsilon^{-1})$. *Then, there exists a choice of* $\{N_\ell\}_{l=0}^L$ *such that*

$$\mathbb{E}_\Delta\left[\mathscr{C}_\varepsilon(\widehat{Q}_{h_L}^{MLQMC})\right] \;=\; \mathscr{O}\Big(\varepsilon^{-2\lambda-\max\left(0,\ \frac{\gamma-\beta\lambda}{\alpha}\right)}\Big)\,. \tag{42}$$

When $\beta\lambda = \gamma$*, then there is an additional factor* $\log(\varepsilon^{-1})^{1+\lambda}$.

This result can be further improved by using higher order QMC rules [13], but we will not consider this here.

It can be shown, for the theoretically optimal values of N_ℓ, that there exists a constant C such that

$$\frac{\mathbb{V}_\Delta[\widehat{Y}_\ell^{QMC}]}{\mathscr{C}_\ell} \;=\; C\,, \tag{43}$$

independently of the level ℓ and of the value of λ (cf. [31, Sect. 3.3]). The same holds for MLMC. This leads to the following adaptive procedure to choose N_ℓ suggested in [20], which we use in our numerical experiments below instead of (37) .

In particular, starting with an initial number of samples on all levels, we alternate the following two steps until $\mathbb{V}[\widehat{Q}_h^{MLMC}] \leq \varepsilon^2/2$:

(i) Estimate $\mathscr{C}_\ell$ and $\mathbb{V}_\Delta[\widehat{Y}_\ell^{QMC}]$ (resp. $\mathbb{V}[\widehat{Y}_\ell^{MC}]$).
(ii) Compute

$$\ell^* = \underset{\ell=0}{\overset{L}{\operatorname{argmax}}} \left(\frac{\mathbb{V}_\Delta[\widehat{Y}_\ell^{QMC}]}{\mathscr{C}_\ell} \right)$$

and double the number of samples on level ℓ^*.

This procedure ensures that, on exit, (43) is roughly satisfied and the numbers of samples across the levels N_ℓ are quasi-optimal.

We use this adaptive procedure for both the MLMC and the MLQMC method. The lack of optimality typically has very little effect on the actual computational cost. Since the optimal formula (37) for MLMC also depends on estimates of $\mathscr{C}_\ell$ and $\mathbb{V}[Y_\ell]$, it sometimes even leads to a better performance. An additional benefit in the case of MLQMC is that the quadrature error in rank-1 lattice rules is typically lowest when the numbers of lattice points is a power of 2.

5 Numerical Results

We now present numerical results to confirm the gains that are possible with the novel multilevel and quasi-Monte Carlo method applied to our 1D NTE model (1), (2), (3). We assume that the scattering cross-section σ_S is a log-normal random field as described in Section 2.1 and that the absorption cross section is constant, $\sigma_A \equiv \exp(0.25)$. We assume no fission, $\sigma_F \equiv 0$, and a constant source term $f = \exp(1)$. We consider two cases, characterised by the choice of smoothness parameter ν in the Matérn covariance function (7). For the first case, we choose $\nu = 0.5$. This corresponds to the exponential covariance and in the following is called the "exponential field". For the second case, denoted the "Matérn field", we choose $\nu = 1.5$. The correlation length and the variance are $\lambda_C = 1$ and $\sigma_{var}^2 = 1$, respectively. The quantity of interest we consider is

$$Q(\omega) = \int_0^1 \phi(x,\omega)\mathrm{d}x \,. \tag{44}$$

For the discretisation, we choose a uniform spatial mesh with mesh width $h = 1/M$ and a quadrature rule (in angle) with $2N = 4M$ points. The KL expansion of $\log(\sigma_S)$ in (8) is truncated after d terms. We choose d sufficiently high, so that the error due to this truncation is negligible while ensuring that the cost of sampling σ_S does not dominate the overall computational cost. In particular, we choose $d = 8h^{-1}$ for the Matérn field and $d = 225h^{-1/2}$ for the exponential field, leading to a maximum of 2048 and 3600 KL modes, respectively, for the finest spatial resolution in each case. We introduce a hierarchy of levels $\ell = 0, ..., L$ corresponding to a sequence of discretisation parameters $h_\ell = 2^{-\ell}h_0$ with $h_0 = 1/4$, and approximate

the quantity of interest in (44) by

$$Q_h(\omega) := \frac{1}{M}\sum_{j=1}^{M} \Phi_{j-1/2}(\omega) .$$

To generate our QMC points we use an (extensible) randomised rank-1 lattice rule (as presented in Section 4.2), with $R = 8$ shifts. We use the generating vector `lattice-32001-1024-1048576.3600`, which is downloaded from [30].

5.1 A Hybrid Direct-Iterative Solver

To compute samples of the neutron flux and thus of the quantity of interest, we propose a hybrid version of the direct and the iterative solver for the Schur complement system (12) described in Section 2.3.

The cost of the iterative solver depends on the number K of iterations that we take. For each ω, we aim to choose K such that the L_2-error $\|\phi(\omega) - \phi^{(K)}(\omega)\|_2$ is less than ε. To estimate K we fix $h = 1/1024$ and $d = 3600$ and use the direct solver to compute ϕ_h for each sample ω. Let $\rho(\omega) := \|\sigma_S(\cdot,\omega)/\sigma(\cdot,\omega)\|_\infty$. For a sufficiently large number of samples, we then evaluate

$$\frac{\log\left(\left\|\phi_h(\omega) - \phi_h^{(K)}(\omega)\right\|_2\right)}{K\log\left(\rho(\omega)\right)}$$

and find that this quotient is less than $\log(0.5)$ in more than 99% of the cases, for $K = 1,\ldots,150$, so that we can choose $C = 0.5$ in (17). We repeat the experiment also for larger values of h and smaller values of d to verify that this bound holds in at least 99% of the cases independently of the discretisation parameter h and of the truncation dimensions d.

Hence, a sufficient, a priori condition to achieve $\|\phi_h(\omega) - \phi_h^{(K)}(\omega)\|_2 < \varepsilon$ in at least 99% of the cases is

$$K = K(\varepsilon,\omega) = \max\left\{1, \left\lceil \frac{\log(2\varepsilon)}{\log\left(\rho(\omega)\right)} \right\rceil \right\}, \tag{45}$$

where $\lceil\cdot\rceil$ denotes the ceiling function. It is important to note that K is no longer a deterministic parameter for the solver (like M or N). Instead, K is a random variable that depends on the particular realisation of σ_S. It follows from (45), using the results in [9, §2], [24] as in Section 3, that $\mathbb{E}[K(\varepsilon,\cdot)] = \mathscr{O}(\log(\varepsilon))$ and $\mathbb{V}[K(\varepsilon,\cdot)] = \mathscr{O}\left(\log(\varepsilon)^2\right)$, with more variability in the case of the exponential field.

Recall from (13) and (15) that, in the case of $N = 2M$, the costs for the direct and iterative solvers are C_1M^3 and C_2KM^2, respectively. In our numerical experiments, we found that in fact $C_1 \approx C_2$, for this particular relationship between M and N. This motivates a third "hybrid" solver, presented in Fig. 1, where the iterative solver is

Given σ_S, σ and desired accuracy ε.

1. Calculate $\rho = \max_x(\sigma_S/\sigma)$.
2. Define $K = \left\lceil \, log(2\varepsilon) \;/\; log(\rho) \, \right\rceil$.
3. If $K < M$, solve using K source iterations. Else, use the direct solver.

Fig. 1 Hybrid direct-iterative solver for the Schur complement system (12) in Section 2.3.

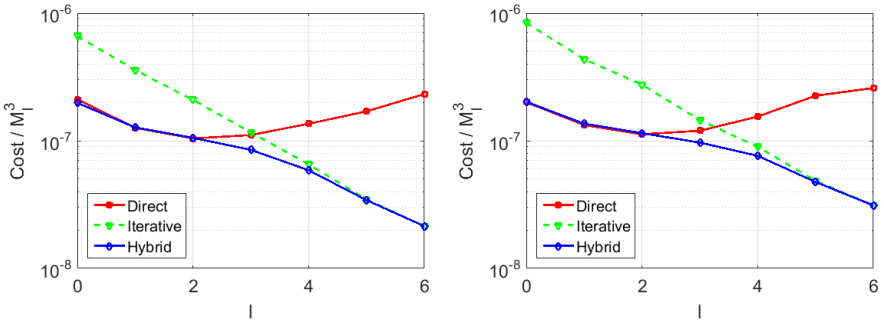

Fig. 2 Comparison of the average costs of the solvers (actual timings in seconds divided by M_ℓ^3) for the Matérn field (left) and for the exponential field (right).

chosen when $K(\omega) < M$ and the direct solver when $K(\omega) \geq M$. By definition, this solver is always faster than the direct or the iterative solver.

We finish this section with a study of actual timings in seconds (here referred to as the cost) of the three solvers. In Fig. 2, we plot the cost divided by M_ℓ^3 and averaged over a sufficiently large number of samples, against the level parameter ℓ. We observe that, as expected, the (scaled) cost of the direct solver is almost constant and the iterative solver is more efficient for larger values of M_ℓ. Over the range of values of M_ℓ considered in our experiments, a best fit for the rate of growth of the cost with respect to the discretisation parameter h_ℓ in (32) is $\gamma \approx 2.2$, for both fields. Thus our solver has a practical complexity of $\mathcal{O}(n^{1.1})$, where $n \sim M^2$ is the total number of degrees of freedom in the system.

5.2 A Priori Error Estimates

Studying the complexity theorems of Section 4, we can see that the effectiveness of the various Monte Carlo methods depends on the parameters α, β, γ and λ in (31), (32), (38) and (41). In this section, we will (numerically) estimate these parameters in order to estimate the theoretical computational cost for each approach.

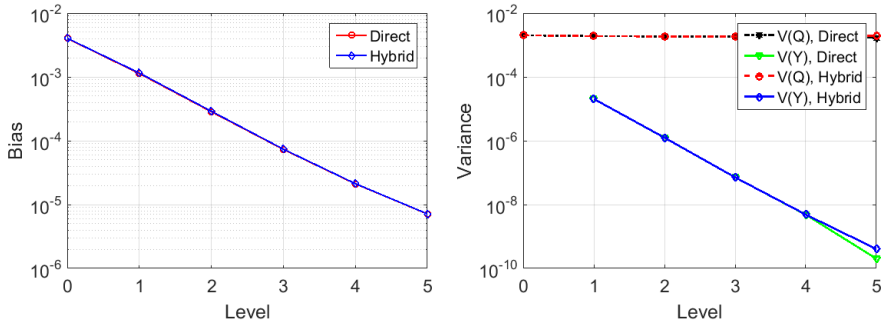

Fig. 3 Estimates of the bias due to discretisation errors (left) and of the variances of Q_{h_ℓ} and Y_ℓ (right), in the case of the Matérn field.

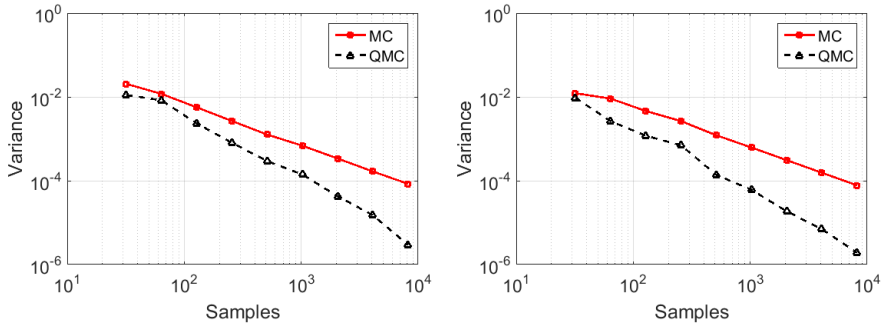

Fig. 4 Convergence of standard Monte Carlo and quasi-Monte Carlo estimators: Matérn field (left) and exponential field (right).

We have already seen that $\gamma \approx 2.2$ for the hybrid solver. In Fig. 3, we present estimates of the bias $\mathbb{E}[Q - Q_{h_\ell}]$, as well as of the variances of Q_{h_ℓ} and of Y_ℓ, computed via sample means and sample variances over a sufficiently large set of samples. We only explicitly show the curves for the Matérn field. The curves for the exponential field look similar. From these plots, we can estimate $\alpha \approx 1.9$ and $\beta \approx 4.1$, for the Matérn field, and $\alpha \approx 1.7$ and $\beta \approx 1.9$, for the exponential field.

To estimate λ in (41), we need to study the convergence rate of the QMC method with respect to the number of samples N_{QMC}. This study is illustrated in Fig. 4. As expected, the variance of the standard MC estimator converges with $\mathscr{O}(N_{MC}^{-1})$. On the other hand, we observe that the variance of the QMC estimator converges approximately with $\mathscr{O}(N_{QMC}^{-1.6})$ and $\mathscr{O}(N_{QMC}^{-1.4})$ (or $\lambda = 0.62$ and $\lambda = 0.71$) for the Matérn field and for the exponential field, respectively.

We summarise all the estimated rates in Table 1.

	α	β	γ	λ
Matérn field	1.9	4.1	2.2	0.62
Exponential field	1.7	1.9	2.2	0.71

Table 1 Summary of estimated rates in (31), (32), (38) and (41).

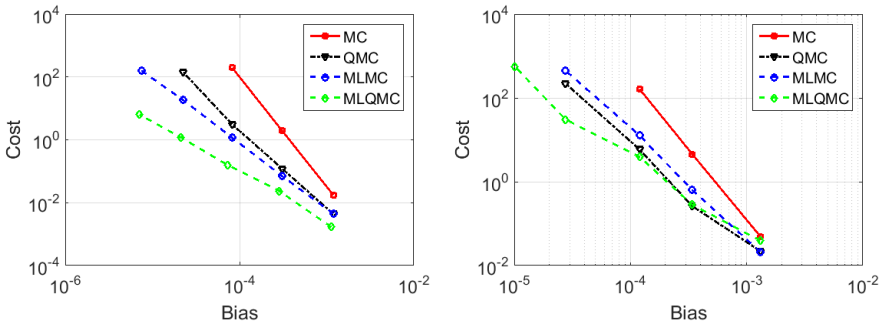

Fig. 5 Actual cost plotted against estimated bias on level L for standard Monte Carlo, QMC, multilevel MC and multilevel QMC: Matérn field (left) and exponential field (right).

5.3 Complexity Comparison of Monte Carlo Variants

For a fair comparison of the complexity of the various Monte Carlo estimators, we now use the a priori bias estimates in Section 5.2 to choose a suitable tolerance ε_L for each choice of $h = h_L$. Let τ_ℓ be the estimated bias on level ℓ. Then, for each $L = 2, \ldots, 6$, we choose $h = h_L$ and $\varepsilon_L := \sqrt{2}\,\tau_L$, and we plot in Fig. 5 the actual cost of each of the estimators described in Section 4 against the estimated bias on level L. The numbers of samples for each of the estimators are chosen such that $\mathbb{V}[\widehat{Q}_h] \le \varepsilon_\ell^2/2$. The coarsest mesh size in the multilevel methods is always $h_0 = 1/4$. We can clearly see the benefits of the QMC sampling rule and of the multilevel variance reduction, and the excellent performance of the multilevel QMC estimator confirms that the two improvements are indeed complementary. As expected, the gains are more pronounced for the smoother (Matérn) field.

We finish by comparing the actual, observed ε-cost of each of the methods with the ε-cost predicted theoretically using the estimates for α, β, γ and λ in Section 5.2. Assuming a growth of the ε-cost proportional to ε^{-r}, for some $r > 0$, we compare in Table 2 estimated and actual rates r for all the estimators. Some of the estimated rates in Section 5.2 are fairly crude, so the good agreement between estimated and actual rates is quite impressive.

	MC		QMC		MLMC		MLQMC	
Field	Estimated	Actual	Estimated	Actual	Estimated	Actual	Estimated	Actual
Matérn	3.2	3.4	2.4	2.7	2.0	2.1	1.2	1.5
Exponential	3.3	3.6	2.7	2.4	2.2	2.5	1.9	1.9

Table 2 Comparison of the estimated theoretical and actual computational ε-cost rates, for different Monte Carlo methods, using the hybrid solver.

6 Conclusions

To summarise, we have presented an overview of novel error estimates for the 1D slab geometry simplification of the Neutron Transport Equation, with spatially varying and random cross-sections. In particular, we consider the discrete ordinates method with Gauss quadrature for the discretisation in angle, and a diamond differencing scheme on a quasi-uniform grid in space. We represent the spatial uncertainties in the cross-sections by log-normal random fields with Matérn covariances, including cases of low smoothness. These error estimates are the first of this kind. They allow us to satisfy key assumptions for the variance reduction in multilevel Monte Carlo methods.

We then use a variety of recent developments in Monte Carlo methods to study the propagation of the uncertainty in the cross-sections, through to a linear functional of the scalar flux. We find that the Multilevel Quasi Monte Carlo method gives us significant gains over the standard Monte Carlo method. These gains can be as large as almost two orders of magnitude in the computational ε-cost for $\varepsilon = 10^{-4}$.

As part of the new developments, we present a hybrid solver, which automatically switches between a direct or iterative method, depending on the rate of convergence of the iterative solver which varies from sample to sample. Numerically, we observe that the hybrid solver is almost an order of magnitude cheaper than the direct solver on the finest mesh, on the other hand the direct solver is almost an order of magnitude cheaper than the iterative solver on the coarsest mesh we considered.

We conclude that modern variants of Monte Carlo based sampling methods are extremely useful for the problem of Uncertainty Quantification in Neutron Transport. This is particularly the case when the random fields are non-smooth and a large number of stochastic variables are required for accurate modelling.

Acknowledgements We thank EPSRC and AMEC Foster Wheeler for financial support for this project and we particularly thank Professor Paul Smith (AMECFW) for many helpful discussions. Matthew Parkinson is supported by the EPSRC Centre for Doctoral Training in Statistical Applied Mathematics at Bath (SAMBa), under project EP/L015684/1. This research made use of the Balena High Performance Computing (HPC) Service at the University of Bath.

References

1. Allen, EJ., Victory Jr, HD., Ganguly, K.: On the convergence of finite-differenced multigroup, discrete-ordinates methods for anisotropically scattered slab media. SIAM J. Numer. Analysis **26**, 88–106 (1989).
2. Asadzadeh, M.: A finite element method for the neutron transport equation in an infinite cylindrical domain. SIAM J. Numer. Analysis **35**, 1299–1314 (1998).
3. Asadzadeh, M., Thevenot, L.: On discontinuous Galerkin and discrete ordinates approximations for neutron transport equation and the critical eigenvalue. Nuovo Cimento C **33**, 21–29 (2010).
4. Ayres, DAF., Eaton, MD.: Uncertainty quantification in nuclear criticality modelling using a high dimensional model representation. Ann. of Nuclear Energy **80**, 379–402 (2015).
5. Ayres, DAF., Park, S., Eaton, MD.: Propagation of input model uncertainties with different marginal distributions using a hybrid polynomial chaos expansion. Ann. of Nuclear Energy **66**, 1–4 (2014).
6. Babuska, I., Nobile, F., Tempone, R.: A stochastic collocation method for elliptic partial differential equations with random input data. SIAM J. Numer. Anal. **45**, 1005-V1034 (2007).
7. Barth, A., Schwab, C., Zollinger, N.: Multi-level Monte Carlo finite element method for elliptic PDEs with stochastic coefficients. Numer. Math. **119**, 123–161 (2011).
8. Blake, J.: Domain decomposition methods for nuclear reactor modelling with diffusion acceleration. PhD Thesis, University of Bath (2016).
9. Charrier, J., Scheichl, R., Teckentrup, AL.: Finite element error analysis of elliptic PDEs with random coefficients and its application to multilevel Monte Carlo methods. SIAM J. Numer. Anal. **51**, 322–352 (2013).
10. Cliffe, KA., Giles, MB., Scheichl, R., Teckentrup, AL.: Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients. Comput. Visual. Sci. **14**, 3–15 (2011).
11. Dautray, R., Lions, JL.: Mathematical Analysis and Numerical Methods for Science and Technology: Volume 1, Physical Origins and Classical Methods. Springer, Heidleberg (2012).
12. Dick, J., Kuo, FY., Le Gia, QT., Nuyens, D., Schwab, C.: Higher order QMC Petrov–Galerkin discretization for affine parametric operator equations with random field inputs. SIAM J. Numer. Anal. **52**, 2676–2702 (2014).
13. Dick, J., Kuo, FY., Le Gia, QT., Schwab, C.: Multi-level higher order QMC Galerkin discretization for affine parametric operator equations. Preprint arXiv:1406.4432, Cornell University (2014).
14. Dick, J., Kuo, FY., Sloan, IH.: High-dimensional integration: the quasi-Monte Carlo way. Acta Num. **22**, 133–288 (2013).
15. Dick, J., Pillichshammer, F.: Digital Nets and Sequences: Discrepancy Theory and QuasiV-Monte Carlo Integration. Cambridge University Press, Cambridge (2010).
16. Eiermann, M., Ernst, O.G., Ullmann, E.: Computational aspects of the stochastic finite element method . Comput. Visual. Sci. **10**, 3–15, (2007).
17. Fichtl, ED., Prinja, AK.: The stochastic collocation method for radiation transport in random media. J. Quantitative Spectroscopy and Radiative Transfer **112**(4), 646–659 (2011).
18. Giles, MB.: Multilevel Monte Carlo path simulation. Oper. Res. **56**, 607–617 (2008).
19. Giles, MB.: Multilevel Monte Carlo methods. Acta Num. **24**, 259–328 (2015).
20. Giles, MB., Waterhouse, BJ.: Multilevel quasi-Monte Carlo path simulation. Adv. Financial Modelling, Radon Series Comp. App. Math., 165–181 (2009).
21. Gilli, L., Lathouwers, D., Kloosterman, JL., van der Hagen, THJJ., Koning, AJ., Rochman, D.: Uncertainty quantification for criticality problems using non-intrusive and adaptive polynomial chaos techniques. Ann. of Nuclear Energy **56**, 71–80 (2013).
22. Goda, T., Dick, J.: Construction of interlaced scrambled polynomial lattice rules of arbitrary high order. Found. Comput. Math. **15**, 1245–1278 (2015).
23. Graham, IG., Kuo, FY., Nuyens, D., Scheichl, R., Sloan, IH.: Quasi-Monte Carlo methods for elliptic PDEs with random coefficients and applications. J. Comput. Phys. **230**, 3668–3694 (2011).

24. Graham, I.G., Kuo, F.Y., Nichols, J.A., Scheichl, R., Schwab, C. and Sloan, I.H.: Quasi-Monte Carlo finite element methods for elliptic PDEs with lognormal random coefficients. Numer. Math. **131**, 329–368 (2015).
25. Graham. IG., Parkinson MJ., Scheichl R.: Error estimates for the neutron transport equation with variable coefficients. In preparation (2017).
26. Gunzburger, M., Webster, CG., Zhang, G.: Stochastic finite element methods for PDEs with random input data. Acta Numer. **23**, 521-V650 (2014).
27. Haji-Ali, AL., Nobile, F., Tempone, R.: Multi-index Monte Carlo: when sparsity meets sampling. Numer. Math. **132**, 767–806 (2016).
28. Heinrich, S.: Multilevel Monte Carlo methods. *Lecture Notes in Computer Science*, vol. 2179, Springer, Heidelberg (2001).
29. Keller, HB.: On the pointwise convergence of the discrete-ordinates method. J. Soc. Indust. Appl. Math. **8**, 560–567 (1960).
30. Kuo, FY.: `http://web.maths.unsw.edu.au/~fkuo/lattice/index.html`
31. Kuo, FY., Scheichl, R., Schwab, C., Sloan, IH., Ullmann, E.: Multilevel quasi-Monte Carlo methods for lognormal diffusion problems. Math. Comput., in press (2016).
32. Kuo, FY., Schwab, C., Sloan, IH.: Quasi-Monte Carlo finite element methods for a class of elliptic partial differential equations with random coefficient. SIAM J. Numer. Anal. **50**, 3351–3374 (2012).
33. Kuo, FY., Schwab, C., Sloan, IH.: Multi-level quasi-Monte Carlo finite element methods for a class of elliptic PDEs with random coefficients. Found. Comput. Math. **15**, 411–449 (2015).
34. Larsen EW., Nelson, P.: Finite difference approximations and superconvergence for the discrete-ordinate equations in slab geometry. SIAM J. Numer. Anal. **19**, 334–348 (1982).
35. Lee, CW., Lee, YO., Cho. YS.: Comparison of the nuclear data libraries in the shielding calculation for the accelerator facility of the Proton Engineering Frontier Project in Korea. Inter. Conf. Nuclear Data for Science and Technology. EDP Sciences, (2007).
36. Lewis, EE., Miller, WF.: Computational methods of Neutron Transport. John Wiley and Sons, New York (1984).
37. Lord, GJ., Powell, CE. Shardlow, T.: An Introduction to Computational Stochastic PDEs. Cambridge University Press, Cambridge (2014).
38. Niederreiter, H.: Quasi-Monte Carlo Methods. John Wiley and Sons, New York (2010).
39. Nobile, F., Tempone, R., Webster, CG.: An anisotropic sparse grid stochastic collocation method for partial differential equations with random input data. SIAM J. Numer. Anal. **46**, 2411V-2442 (2008).
40. Nuyens, D., Cools, R.: Fast algorithms for component-by-component construction of rank-1 lattice rules in shift-invariant reproducing kernel Hilbert spaces. Math. Comput. **75**, 903–920 (2006).
41. Pitkaranta, J., Scott, LR.: Error estimates for the combined spatial and angular approximations of the transport equation for slab geometry. SIAM J. Numer. Anal. **20**, 922–950 (1983).
42. Reed, WH., Hill, TR.: Triangular mesh methods for the neutron transport equation. Technical Report LA-UR-73-479, Los Alamos National Laboratory (1973).
43. Sanchez, R., McCormick, NJ.: Review of neutron transport approximations. Nucl. Sci. Eng. US **80**, 481V-535 (1982).
44. Sloan, IH.: Error analysis for a class of degenerate-kernel methods, Numer. Math. **25**, 231-238 (1975).
45. Sloan, IH., Wozniakowski, H.: When are quasi-Monte Carlo algorithms efficient for high dimensional integrals? J. Complexity **14**, 1–33 (1998).
46. Teckentrup, AL., Jantsch, P., Webster, CG., Gunzburger, M.: A multilevel stochastic collocation method for partial differential equations with random input data. SIAM/ASA J Uncertainty Quant. **3**, 1046–1074 (2015).
47. Victory Jr, HD.: Convergence of the multigroup approximations for subcritical slab media and applications to shielding calculations. Adv. Appl. Math. **5**, 227–259 (1984).
48. Xiu, D., Karniadakis, GE.: The Wiener-Askey polynomial chaos for stochastic differential equations. SIAM J. Sci. Comput. **24**, 614–644 (2002).